

From Closed-world Enforcement to Open-world Assessment of Privacy

Michael Backes^{1,2}

Pascal Berrang¹

Praveen Manoharan¹

¹CISPA, Saarland University ² MPI-SWS

backes@cs.uni-saarland.de

berrang@cs.uni-saarland.de

manoharan@cs.uni-saarland.de

Abstract

In this paper, we develop a *user-centric privacy framework for quantitatively assessing the exposure* of personal information in open settings. Our formalization addresses key-challenges posed by such open settings, such as the unstructured dissemination of heterogeneous information and the necessity of user- and context-dependent privacy requirements. We propose a new definition of information sensitivity derived from our formalization of privacy requirements, and, as a sanity check, show that hard non-disclosure guarantees are impossible to achieve in open settings.

After that, we provide an instantiation of our framework to address the *identity disclosure* problem, leading to the novel notion of *d-convergence*. *d-convergence* is based on indistinguishability of entities and it bounds the likelihood with which an adversary successfully links two profiles of the same user across online communities.

Finally, we provide a *large-scale evaluation* of our framework on a collection of 15 million comments collected from the Online Social Network Reddit. Our evaluation validates the notion of *d-convergence* for assessing the linkability of entities in our data set and provides deeper insights into the data set's structure.

Contents

1	Introduction	3
1.1	Contribution	4
1.2	Outline	5
2	Related Work	5
3	Privacy in Open Settings	7
3.1	Example	7
3.2	Challenges of Privacy in Open Settings	8
3.3	Inadequacy of Existing Models	9
4	A Framework for Privacy in Open Settings	10
4.1	Modeling Information in Open Settings	10
4.2	Adversary Model	11
4.3	Inapplicability of Statistical Privacy Notions	12
4.4	User-Specified Privacy Requirements	14
4.5	Sensitive Information	15
5	Linkability in Open Settings	16
5.1	Model Instantiation for Linkability	16
5.2	Anonymity	17
5.3	Entity Matching	17
5.4	Identity Disclosure	18
5.5	Limitations	19
6	Linkability Evaluation on Reddit	20
6.1	Goals	20
6.2	Data-Collection	20
6.3	Model Instantiation	21
6.4	Data-Processing	22
6.5	Evaluation and Discussion	25
7	Conclusion and Future Work	27

1 Introduction

The Internet has undergone dramatic changes in the last two decades, evolving from a mere communication network to a global multimedia platform in which billions of users not only actively exchange information, but increasingly conduct sizable parts of their daily lives. While this transformation has brought tremendous benefits to society, it has also created new threats to online privacy that existing technology is failing to keep pace with. Users tend to reveal personal information without considering the widespread, easy accessibility, potential linkage and permanent nature of online data. Many cases reported in the press show the resulting risks, which range from public embarrassment and loss of prospective opportunities (e.g., when applying for jobs or insurance), to personal safety and property risks (e.g., when sexual offenders or burglars learn users’ whereabouts online). The resulting privacy awareness and privacy concerns of Internet users have been further amplified by the advent of the Big-Data paradigm and the aligned business models of personalized tracking and monetizing personal information in an unprecedented manner.

Developing a suitable methodology to reason about the privacy of users in such a large-scale, open web setting, as well as corresponding tool support in the next step, requires at its core a formal privacy model that lives up to the now increasingly dynamic dissemination of unstructured, heterogeneous user content on the Internet: While users traditionally shared information mostly using public profiles with static information about themselves, nowadays they disseminate personal information in an unstructured, highly dynamic manner, through content they create and share (such as blog entries, user comments, a “Like” on Facebook), or through the people they befriend or follow. Furthermore, ubiquitously available background knowledge about a dedicated user needs to be appropriately reflected within the model and its reasoning tasks, as it can decrease a user’s privacy by inferring further sensitive information. As an example, Machine Learning and other Information Retrieval techniques provide comprehensive approaches for profiling a user’s actions across multiple Online Social Networks, up to a unique identification of a given user’s profiles for each such network.

Prior research on privacy has traditionally focused on closed database settings – characterized by a complete view on structured data and a clear distinction of key- and sensitive attributes – and has aimed for strong privacy guarantees using global data sanitization. These approaches, however, are inherently inadequate if such closed settings are replaced by open settings as described above, where unstructured and heterogeneous data is being disseminated, where individuals have a partial view of the available information, and where global data sanitization is impossible and hence strong guarantees have to be replaced by probabilistic privacy assessments.

As of now, *even the basic methodology is missing* for offering users technical means to comprehensively assess the privacy risks incurred by their data dissemination, and their daily online activities in general. Existing privacy models such as k -anonymity [30], l -diversity [22], t -closeness [21] and the currently most

popular notion of Differential Privacy [11] follow a database-centric approach that is inadequate to meet the requirements outlined above. We refer the reader to Section 3.3 for further discussions on existing privacy models.

1.1 Contribution

In this paper, we present a rigorous methodology for quantitatively assessing the exposure of personal information in open settings. Concretely, the paper makes the following three tangible contributions: (1) a formal framework for reasoning about the disclosure of personal information in open settings, (2) an instantiation of the framework for reasoning about the identity disclosure problem, and (3) an evaluation of the framework on a collection of 15 million comments collected from the Online Social Network Reddit.

A Formal Framework for Privacy in Open Settings. We propose a novel framework for addressing the essential challenges of privacy in open settings, such as providing a data model that is suited for dealing with unstructured dissemination of heterogeneous information through various different sources and a flexible definition of user-specific privacy requirements that allow for the specification of context-dependent privacy goals. In contrast to most existing approaches, our framework strives to assess the degree of exposure individuals face, in contrast to trying to enforce an individual’s privacy requirements. Moreover, our framework technically does not differentiate between non-sensitive and sensitive attributes a-priori, but rather starts from the assumption that all data is equally important and can lead to privacy risks. More specifically, our model captures the fact that the sensitivity of attributes is highly user- and context-dependent by deriving information sensitivity from each user’s privacy requirements. As a sanity check we prove that hard non-disclosure guarantees cannot be provided for the open setting in general, providing incentive for novel approaches for assessing privacy risks in the open settings.

Reasoning about Identity Disclosure in Open Settings. We then instantiate our general privacy framework for the specific use case of identity disclosure. Our framework defines and assesses identity disclosure (i.e., identifiability and linkability of identities) by utilizing entity similarity, i.e., an entity is private in a collection of entities if it is sufficiently similar to its peers. At the technical core of our model is the new notion of d -convergence, which quantifies the similarity of entities within a larger group of entities. It hence provides the formal grounds to assess the ability of any single entity to blend into the crowd, i.e., to hide amongst peers. The d -convergence model is furthermore capable of assessing identity disclosure risks specifically for single entities. To this end, we extend the notion of d -convergence to the novel notion of (k, d) -anonymity, which allows for entity-centric identity disclosure risk assessments by requiring d -convergence in the local neighborhood of a given entity. Intuitively, this new notion provides a generalization of k -anonymity that is not bound to matching identities based on pre-defined key-identifiers.

Empirical Evaluation on Reddit. Third, we perform an instantiation of our identity disclosure model for the important use case of analyzing user-generated text content in order to characterize specific user profiles. We use unigram frequencies extracted from user-generated content as user attributes, and we subsequently demonstrate that the resulting unigram model can indeed be used for quantifying the degree of anonymity of – and ultimately, for differentiating – individual entities. For the sake of exposition, we apply this unigram model to a collection of 15 million comments collected from the Online Social Network Reddit. The computations were performed on two Dell PowerEdge R820 with 64 virtual cores each at 2.60GHz over the course of six weeks. Our evaluation shows that (k, d) -anonymity suitably assesses an identity’s anonymity and provides deeper insights into the data set’s structure.

1.2 Outline

We begin by discussing related work in Section 2 and explain why existing privacy notions are inadequate for reasoning about privacy in open web settings in Section 3. We then define our privacy framework in Section 4 and instantiate it for reasoning about identity disclosure in Section 5. In Section 6 we perform a basic evaluation of the identity disclosure model on the Reddit Online Social Network. We summarize our findings 7.

2 Related Work

In this section, we give an overview over other relevant related work that has not yet been considered in the previous subsection.

Privacy in Closed-world Settings. The notion of privacy has been exhaustively discussed for specific settings such as statistical databases, as well as for more general settings. Since we already discussed the notions of k -anonymity [30], l -diversity [22] t -closeness [21] and Differential Privacy [11] in Section 3.3 in great detail, we will now discuss further such notions.

A major point of criticism of Differential Privacy, but also the other existing privacy notions, found in the literature [18] is the (often unclear) trade-off between utility and privacy that is incurred by applying database sanitation techniques to achieve privacy. Several works have shown that protection against attribute disclosure cannot be provided in settings that consider an adversary with arbitrary auxiliary information [10, 12, 13]. We later show, as sanity check, that in our formalization of privacy in open settings, general non-disclosure guarantees are indeed impossible to achieve. By providing the necessary formal groundwork in this paper, we hope to stimulate research on *assessing* privacy risks in open settings, against explicitly spelled-out adversary models.

Kasiviswanathan and Smith [17] define the notion of ϵ -semantic privacy to capture general non-disclosure guarantees. We define our adversary model in a similar fashion as in their formalization and we use ϵ -semantic privacy to show

that general non-disclosure guarantees cannot be meaningfully provided in open settings.

Several extensions of the above privacy notions have been proposed in the literature to provide privacy guarantees in use cases that differ from traditional database privacy [3, 6, 34, 16, 36, 7]. These works aim at suitably transforming different settings into a database-like setting that can be analyzed using differential privacy. Such a transformation, however, often abstracts away from essential components of these settings, and as a result achieve impractical privacy guarantees. As explained in Section 3.3, the open web setting is particularly ill-suited for such transformations.

Specifically for the use case in Online Social Networks (in short, OSNs), many works [34, 19, 16, 36, 7] apply the existing database privacy notions for reasoning about attribute disclosure in OSN data. These works generally impose a specific structure on OSN data, such as a social link graph, and reason about the disclosure of private attributes through this structure. Zhaleva et al. [34] show that mixed public and private profiles do not necessarily protect the private part of a profile since they can be inferred from the public part. Heatherly et al. [16] show how machine learning techniques can be used to infer private information from publicly available information. Kosinski et al. [19] moreover show that machine learning techniques can indeed be used to predict personality traits of users and their online behavior. Zhou et al. [36] apply the notions of k -anonymity and l -diversity to data protection in OSNs and discuss the complexity of finding private subsets. Their approach does however suffer from the same problems these techniques have in traditional statistical data disclosure, where an adversary with auxiliary information can easily infer information about any specific user. Chen et al. [7] provide a variation of differential privacy which allows for privacy and protection against edge-disclosure attacks in the correlated setting of OSNs. The setting, however, remains static, and it is assumed that the data can be globally sanitized in order to provide protection against attribute disclosure. Again, as discussed in Section 3.3, this does not apply to the open web setting with its highly unstructured dissemination of data.

Statistical Language Models. Statistical Language Models for information retrieval have first been introduced by Ponte and Croft [26] as an alternative approach for document retrieval and are inspired by language models for Speech Recognition and Natural Language Processing [29, 27]. They have subsequently been focus of a long line of research (examples include [20, 33, 31, 32]) that further develop the basic statistical language model approach and its benefits. While Statistical Language Models have not been shown to perform better than other established retrieval methods [32], we found that the Statistical Language Model formulation is closer than other options to what we require in expressing and solving indistinguishability problems that arise in computer security.

3 Privacy in Open Settings

Before we delve into the technical parts of this paper, we give an informal overview over privacy in the Internet of the future. To this end, we first provide an example that illustrates some of the aspects of privacy in the Internet, and then in detail discuss the challenges of privacy in the Internet and why existing privacy notions are not applicable to this setting.

3.1 Example

Consider the following example: Employer Alice receives an application by potential employee Bob which contains personal information about Bob. Before she makes the decision on the employment of Bob, however, she searches the internet and tries to learn even more about her potential employee. A prime source of information are, for example, Online Social Networks (OSNs) which Alice can browse through. If she manages to identify Bob's profile in such an OSN she can then learn more about Bob by examining the publicly available information of this profile.

In order to correctly identify Bob's profile in an OSN, Alice takes the following approach: based on the information found in Bob's application, she constructs a model θ_B that contains all attributes, such as name, education or job history, extracted from Bob's application. She then compares this model θ_B to the profiles P_1, \dots, P_n found in the OSNs and ranks them by similarity to the model θ_B . Profiles that show sufficient similarity to the model θ_B are then chosen by Alice as belonging to Bob. After identifying the (for Alice) correctly matching profiles P_1^*, \dots, P_i^* of Bob, Alice can finally merge their models $\theta_1^*, \dots, \theta_i^*$ with θ_B to increase her knowledge about Bob.

Bob now faces the problem that Alice could learn information about him that he does not want her to learn. He basically has two options: he either does not share this critical information at all, or makes sure that his profile is not identifiable as his. In OSNs such as Facebook, where users are required to identify themselves, Bob can only use the first option. In anonymous or pseudonymous OSNs such as Reddit or Twitter, however, he can make use of the second option. He then has to make sure that he does not share enough information on his pseudonymous profiles that would allow Alice to link his pseudonymous profile to him personally.

Privacy in the open web is mostly concerned with the second option: we cannot protect an entity ϵ against sharing personal information through a profile which is already uniquely identified with the entity ϵ . We can, however, estimate how well an pseudonymous account of ϵ can be linked to ϵ , and through this link, learn personal information about ϵ . As the example above shows, we can essentially measure privacy in terms of similarity of an entity ϵ in a collection of entities \mathcal{E} .

The identifiability of ϵ then substantially depends on the attributes ϵ exhibits in the context of \mathcal{E} and does not necessarily follow the concept of personally identifiable information (PII) as known in the more common understanding of pri-

vacy and in privacy and data-protection legislation [2]: here, privacy protection only goes as far as protecting this so-called personally identifiable information, which often is either not exactly defined, or restricted to an a-priori-defined set of attributes such as name, Social Security number, etc. We, along with other authors in the literature [25, 24], find however that the set of critical attributes that need to be protected differ from entity to entity, and from community to community. For example, in a community in which all entities have the name “Bob”, exposing your name does not expose any information about yourself. In a different community, however, where everyone has a different name, exposing your name exposes a lot of information about yourself.

In terms of the privacy taxonomy formulated by Zheleva and Getoor [35], the problem we face corresponds to the identity disclosure problem, where one tries to identify whether and how an identity is represented in an OSN. We think that this is one of the main concerns of users of frequently used OSNs, in particular those that allow for pseudonymous interactions: users are able to freely express their opinions in these environments, assuming that their opinions cannot be connected to their real identity. However, any piece of information they share in their interactions can leak personal information that can lead to identity disclosure, defeating the purpose of such pseudonymous services.

To successfully reason about the potential disclosure of sensitive information in such open settings, we first have to consider various challenges that have not been considered in traditional privacy research. After presenting these challenges, we discuss the implications of these challenges on some of the existing privacy notions, before we consider other relevant related work in the field.

3.2 Challenges of Privacy in Open Settings

In this subsection, we introduce the challenges induced by talking about privacy in open settings:

C1) Modeling heterogeneous information. We require an information model that allows for modeling various types of information and that reflects the heterogeneous information shared throughout the Internet. This model needs to adequately represent personal information that can be inferred from various sources, such as static profile information or from user-generated content, and should allow statistical assessments about the user, as is usually provided by knowledge inference engines. We propose a solution to this challenge in Section 4.1.

C2) User-specified privacy requirements. We have to be able to formalize user-specified privacy requirements. This formalization should use the previously mentioned information model to be able to cope with heterogeneous information, and specify which information should be protected from being publicly disseminated. We present a formalization of user privacy requirements in Section 4.4.

C3) Information sensitivity. In open settings, information sensitivity is a function of user expectations and context: we therefore need to provide new

definitions for sensitive information that takes user privacy requirements into account. We present context- and user-specific definitions of information sensitivity in Section 4.5.

C4) Adversarial knowledge estimation. To adequately reason about disclosure risks in open settings we also require a parameterized adversary model that we can instantiate with various assumptions on the adversary’s knowledge: this knowledge should include the information disseminated by the user, as well as background knowledge to infer additional information about the user. In Section 4, we define our adversary model based on statistical inference.

In the following sections, we provide a rigorous formalization for these requirements, leading to a formal framework for privacy in open settings. We will instantiate this framework in Section 5.3 to reason about the identity disclosure in particular.

We begin by discussing why existing privacy notions are not suited for reasoning about privacy in open settings. Afterwards, we provide an overview over further related work.

3.3 Inadequacy of Existing Models

Common existing privacy notions such as k -anonymity [30], l -diversity [22], t -closeness [21] and the currently most popular notion of Differential Privacy [11] provide the technical means for privacy-friendly data-publishing in a closed-world setting: They target scenarios in which all data is available from the beginning, from a single data source, remains static and is globally sanitized in order to provide rigorous privacy guarantees. In what follows, we describe how these notions fail to adequately address the challenges of privacy in open settings discussed above.

a) Absence of structure and classification of data. All the aforementioned privacy models require an a-priori structure and classification of the data under consideration. Any information gathered about an individual thus has to be embedded in this structure, or it cannot be seamlessly integrated in these models.

b) No differentiation of attributes. All of these models except for Differential Privacy require an additional differentiation between key attributes that identify an individual record, and sensitive attributes that a users seeks to protect. This again contradicts the absence of an a-priori, static structure in our setting. Moreover, as pointed out above and in the literature [25], such a differentiation cannot be made a-priori in general, and it would be highly context-sensitive in the open web setting.

c) Ubiquitously available background knowledge. All of these models, except for Differential Privacy, do not take into account adversaries that utilize ubiquitously available background knowledge about a target user to infer additional sensitive information. A common example of background knowledge is openly

available statistical information that allows the adversary to infer additional information about an identity.

d) Privacy for individual users. All these models provide privacy for the whole dataset, which clearly implies privacy of every single user. One of the major challenges in open settings such as the Internet, however, is that accessing and sanitizing all available data is impossible. This leads to the requirement to design a local privacy notion that provides a lower privacy bound for every individual user, even if we only have partial access to the available data.

The notion of Differential Privacy only fails to address some of the aforementioned requirements (parts *a* and *d*), but it comes with the additional assumption that the adversary knows almost everything about the data set in question (everything except for the information in one database entry). This assumption enables Differential Privacy to avoid differentiation between key attributes and sensitive attributes. This strong adversarial model, however, implies that privacy guarantees are only achievable if the considered data is globally perturbed [10, 12, 13], which is not possible in open web settings.

The conceptual reason for the inadequacy of existing models for reasoning about privacy in open web settings is mostly their design goal: Privacy models have thus far mainly been concerned with the problem of attribute disclosure within a single data source: protection against identity disclosure was then attempted by preventing the disclosure of any (sensitive) attributes of a user to the public. In contrast to static settings such as private data publishing, where we can decide which information will be disclosed to the adversary, protection against any attribute disclosure in open settings creates a very different set of challenges which we will address in the following sections.

4 A Framework for Privacy in Open Settings

In this section, we first develop a user model that is suited for dealing with the information dissemination behavior commonly observed on the Internet. We then formalize our adversary model and show, as a sanity check, that hard privacy guarantees cannot be achieved in open settings. We conclude by defining privacy goals in open settings through user-specified privacy requirements from which we then derive a new definition of information sensitivity suited to open settings.

4.1 Modeling Information in Open Settings

We first define the notion of entity models and restricted entity models. These models capture the behavior of these entities and in particular describe which attributes an entity exhibits publicly.

Definition 1 (Entity Model). *Let \mathcal{A} be the set of all attributes. The entity model θ_ϵ of an entity ϵ provides for all attributes $\alpha \in \mathcal{A}$ an attribute value*

$\theta_\epsilon(\alpha) \in \text{dom}(\alpha) \cup \{\text{NULL}\}$ where $\text{dom}(\alpha)$ is the domain over which the attribute α_i is defined.

The domain $\text{dom}(\theta)$ of an entity model θ is the set of all attributes $\alpha \in \mathcal{A}$ with value $\theta(\alpha) \neq \text{NULL}$.

An entity model thus corresponds to the information an entity can publicly disseminate. With the specific null value **NULL** we can also capture those cases where the entity does not have any value for that specific attribute.

In case the adversary has access to the full entity model, a set of entity models basically corresponds to a database with each attribute $\alpha \in \mathcal{A}$ as its columns. In the open setting, however, an entity typically does not disseminate all attribute values, but instead only a small part of them. We capture this with the notion of restricted entity models.

Definition 2 (Restricted Entity Model). *The restricted entity model $\theta_\epsilon^{\mathcal{A}'}$ is the entity model of ϵ restricted to the non empty attribute set $\mathcal{A}' \neq \emptyset$, i.e.,*

$$\theta_\epsilon^{\mathcal{A}'}(\alpha) = \begin{cases} \theta_\epsilon(\alpha), & \text{if } \alpha \in \mathcal{A}' \\ \text{NULL}, & \text{otherwise} \end{cases}$$

In the online setting, each of the entities above corresponds to an online profile. A user u usually uses more than one online service, each with different profiles P_1^u, \dots, P_l^u . We thus define a user model as the collection of the entity models describing each of these profiles.

Definition 3 (User Model / Profile Model). *The user model $\theta_u = \{\theta_{P_1^u}, \dots, \theta_{P_l^u}\}$ of a user u is a set of the entity models $\theta_{P_1^u}, \dots, \theta_{P_l^u}$, which we also call profile models.*

With a user model that separates the information disseminated under different profiles, we will be able to formulate privacy requirements for each of these profiles separately. We will investigate this in Section 4.4.

4.2 Adversary Model

In the following we formalize the adversary we consider for privacy in open settings. In our formalization, we follow the definitions of a semantic, Bayesian adversary introduced by Kasiviswanathan and Smith [17].

For any profile P , we are interested in what the adversary **Adv** learns about P observing publicly available information from P . We formalize this learning process through *beliefs* on the models of each profile.

Definition 4 (Belief). *Let \mathcal{P} be the set of all profiles and let $\mathcal{D}_\mathcal{A}$ be the set of all distributions over profile models. A belief $b = \{b_P | P \in \mathcal{P}\}$ is a set of distributions $b_P \in \mathcal{D}_\mathcal{A}$.*

We can now define our privacy adversary in open settings using the notion of belief above.

Definition 5 (Adversary). *An adversary Adv is a pair of prior belief b and world knowledge κ , i.e., $\text{Adv} = (b, \kappa)$.*

The adversary Adv 's prior belief b represents his belief in each profile's profile model before makes any observations. This prior belief can, in particular, also include background knowledge about each profile P . The world knowledge κ of the adversary represents a set of inference rules that allow him to infer additional attribute values about each profile from his observations.

We next define the publicly observations based on which the adversary learns additional information about each profile.

Definition 6 (Publication Function). *A publication function G is a randomized function that maps each profile model θ_P to a restricted profile model $G(\theta_P) = \theta_P^{A'}$ such that there is at least one attribute $\alpha \in \mathcal{A}'$ with $\theta_P(\alpha) = G(\theta_P)(\alpha)$.*

The publication function G reflects which attributes are disseminated publicly by the user through his profile P . G can, in particular, also include local sanitization where some attribute values are perturbed. However, we do require that at least one attribute value remains correct to capture utility requirements faced in open settings.

A public observation now is the collection of all restricted profile models generated by a publication function.

Definition 7 (Public Observation). *Let \mathcal{P} be the set of all profiles, and let G be a publication function. The public observation \mathcal{O} is the set of all restricted profile models generated by G , i.e., $\mathcal{O} = \{G(\theta_P) | P \in \mathcal{P}\}$.*

The public observation \mathcal{O} essentially captures all publicly disseminated attribute values that can be observed by the adversary. Given such an observation \mathcal{O} , we can now determine what the adversary Adv learns about each profile by determining his *a-posteriori belief*.

Definition 8 (A-Posteriori Belief). *Let \mathcal{P} be the set of all profiles. Given an adversary $\text{Adv} = (b, \kappa)$ and a public observation \mathcal{O} , the adversary's a-posteriori belief $\bar{b} = \{\bar{b}_P \in \mathcal{D}_A | P \in \mathcal{P}\}$ is determined by applying the Bayesian inference rule, i.e.,*

$$\bar{b}_P[\theta | \mathcal{O}, \kappa] = \frac{Pr[\mathcal{O} | \kappa, \theta] \cdot b_P[\theta]}{\sum_{\theta'} Pr[\mathcal{O} | \kappa, \theta'] \cdot b_P[\theta']}.$$

Here, the conditional probability $Pr[\mathcal{O} | \kappa, \theta]$ describes the likelihood that the observational \mathcal{O} is created by the specific entity model θ .

We will utilize the a-posteriori belief of the adversary to reason about the violation of the user specified privacy requirements in Section 4.4.

4.3 Inapplicability of Statistical Privacy Notions

In the following, we formally show that traditional non-disclosure guarantees, e.g., in the style of Differential Privacy, are not possible in open settings.

Kasiviswanathan and Smith [17] provide a general definition of non-disclosure they call ϵ -privacy. In their definition, they compare the adversary Adv 's a-posteriori beliefs after observing the transcript t generated from a database sanitization mechanism \mathcal{F} applied on two adjacent databases with n rows: first on the database x , leading to the belief $\bar{b}_0[.]|t$, and secondly on the database x_{-i} , where a value in the i th row in x is replaced by a default value, leading to the belief $\bar{b}_i[.]|t$.

Definition 9 (ϵ -semantic Privacy [17]). *Let $\epsilon \in [0, 1]$. A randomized algorithm \mathcal{F} is ϵ -semantically private if for all belief distributions b on D^n , for all possible transcripts, and for all $i = 1 \dots n$:*

$$\text{SD}(\bar{b}_0[.]|t, \bar{b}_i[.]|t) \leq \epsilon.$$

Here, SD is the total variation distance of two probability distributions.

Definition 10. *Let X and Y be two probability distributions over the sample space D . The total variation distance SD of X and Y is*

$$\text{SD}(X, Y) = \max_{S \subseteq D} [Pr[X \in S] - Pr[Y \in S]].$$

Kasiviswanathan and Smith [17] show that ϵ -differential privacy is essentially equivalent to ϵ -semantic privacy.

In our formalization of privacy in open settings, varying a single database entry corresponds to changing the value of a single attribute α in the profile model θ_P of a profile P to a default value. We denote this modified entity model with θ_P^α , and the thereby produced a-posteriori belief by \bar{b}_P^α . A profile P would then be ϵ -semantically private if for any modified profile model θ_P^α , the a-posteriori belief of adversary Adv does not change by more than ϵ .

Definition 11 (ϵ -semantic Privacy in Open Settings). *Let $\epsilon \in [0, 1]$. A profile P is ϵ -semantically private in open settings if for any attribute α ,*

$$\text{SD}(\bar{b}_P[.]|\mathcal{O}, \bar{b}_P^\alpha[.]|\mathcal{O}) \leq \epsilon$$

where \bar{b}_P and \bar{b}_P^α are the a-posteriori beliefs of the adversary after observing the public output of θ_P and θ_P^α respectively.

As expected, we can show that ϵ -semantic privacy can only hold for $\epsilon = 1$ in open settings.

Theorem 1. *For any profile model θ_P and any attribute α , there is an adversary Adv such that*

$$\text{SD}(\bar{b}[.]|\mathcal{O}, \bar{b}^\alpha[.]|\mathcal{O}) \geq 1.$$

Proof. Let Adv have a uniform prior belief, i.e., all possible profile models have the same probability, and empty world knowledge κ . Let α be the one attribute that remains the same after applying the publication function G . Let x be the original value of this attribute α and let x^* be the default value that replaces x .

Observing the restricted profile model $\theta_P[\mathcal{A}']$ without any additional world knowledge will lead to an a-posteriori belief, where the probability of the entity model θ with $\theta[\mathcal{A}'] = \theta_P[\mathcal{A}']$ and NULL everywhere else, is set to 1.

Conversely, the modified setting will result in an a-posteriori belief that sets the probability for the entity model θ^* to one, where θ^* is constructed for the modified setting as θ above. Thus $\bar{b}[\theta|\mathcal{O}] = 1$, whereas $\bar{b}^\alpha[\theta|\mathcal{O}] = 0$, and hence $SD(\bar{b}[\cdot|\mathcal{O}], \bar{b}^\alpha[\cdot|\mathcal{O}]) = 1$. \square

Intuitively, the adversary can easily distinguish differing profile models because a) he can directly observe the profiles publicly available information, b) he chooses which attributes he considers for his inference and c) only restricted, local sanitization is available to the profile. Since these are elementary properties of privacy in open settings, we can conclude that hard security guarantees in the style of differential privacy are impossible to achieve in open settings.

However, we can provide an assessment of the disclosure risks by explicitly fixing the a-priori knowledge and the attribute set considered by the adversary. While we no longer all-quantify over all possible adversaries, and therefore lose the full generality of traditional non-disclosure guarantees, we might still provide meaningful privacy assessments in practice. We further discuss this approach in Section 4.5, and follow this approach in our instantiation of the general model for assessing the likelihood of identity disclosure in Section 5.

4.4 User-Specified Privacy Requirements

In the following we introduce user-specified privacy requirements that allow us to formulate privacy goals that are user- and context-dependent. These can then lead to restricted privacy assessments instead of general privacy guarantees that we have shown to be impossible in open setting in the previous section.

We define a user's privacy requirements on a per-profile basis, stating which attribute values should not be inferred by adversary after seeing a public observations \mathcal{O} .

Definition 12 (Privacy Policy). *A privacy policy \mathcal{R} is a set of privacy requirements $r = (P, \{\alpha_i = x_i\})$ which require that profile P should never expose the attribute values x_i for the attributes $\alpha_i \in \mathcal{A}$.*

By setting privacy requirements in a per-profile basis we capture an important property of information dissemination in open settings: users utilize different profiles for different context (e.g., different online services) assuming these profiles remain separate and specific information is only disseminated under specific circumstances.

Given the definition of privacy policies, we now define the violation of a policy by considering the adversary's a-posterior belief \bar{b} , as introduced in Section 4.2.

Definition 13 (Privacy Policy Satisfaction / Violation). *Let $\text{Adv} = (b, \kappa)$ be an adversary with a-posteriori belief \bar{b} , and let $\theta[\alpha = x]$ be the set of all entity models that have the value x for the attribute α . A profile P_i^u σ -satisfies a user's privacy requirement $r_j^u = (P, \{\alpha_i = x_i\})$, written $P_i^u \models_\sigma r_j^u$, if*

- $P = P_i^u$
- $\forall \alpha_i : \sum_{\theta \in \theta[\alpha_i = x_i]} \bar{b}_P[\theta | \mathcal{O}, \kappa] \leq \sigma$

and σ -violates the user's privacy requirement otherwise.

A user model θ_u σ -satisfies a user u 's privacy policy \mathcal{R}_u , written $\theta_u \models_\sigma \mathcal{R}_u$, if all profile models $\theta_{P_i^u}$ σ -satisfy their corresponding privacy requirements, and σ -violates the privacy policy otherwise.

The above attributes can also take the form of “ P belongs to the same user as P' ”, effectively restricting which profiles should be linked to each other. We will investigate this profile linkability problem specifically in Section 5.

4.5 Sensitive Information

In contrast to the closed-world setting, with its predefined set of sensitive attributes that automatically defines the privacy requirements, a suitable definition of information sensitivity in the open setting is still missing. In the following, we derive the notion of sensitive information from the user privacy requirements we defined in Section 4.4.

Definition 14 (Sensitive Attributes). *A set of attributes \mathcal{A}^* is sensitive for a user u in the context of her profile P_i^u if u 's privacy policy \mathcal{R}_u contains a privacy requirement $r = (P_i^u, \mathcal{A}' = X)$ where $\mathcal{A}^* \subseteq \mathcal{A}'$.*

Here, we use the notation $\mathcal{A} = X$ as vector representation for $\forall \alpha_i \in \mathcal{A} : \alpha_i = x_i$.

Sensitive attributes, as defined above, are not the only type of attributes that are worth to protect: In practice, an adversary can additionally infer sensitive attributes from other attributes through statistical inference using a-priori knowledge. We call such attributes that allow for the inference of sensitive attributes *critical attributes*.

Definition 15 (Critical Attributes). *Given a set of attributes \mathcal{A}^* , let P be a profile with $\text{dom}(\theta_P) \supseteq \mathcal{A}$, and let P' be the profile with the restricted profile model $\theta_{P'} = \theta_P^{\mathcal{A}'}$, where $\mathcal{A}' = \text{dom}(\theta_P) \setminus \mathcal{A}^*$.*

The set of attributes \mathcal{A}^ is σ -critical for the user u that owns the profile P and an adversary with prior belief b_P and world knowledge κ , if u 's privacy policy \mathcal{R}_u contains a privacy requirement r such that P σ -violates r but P' does not.*

Critical information require the same amount of protection as sensitive information, the difference however being that critical information is only protected for the sake of protecting sensitive information.

As a direct consequence of the definition above, sensitive attributes are also critical.

Corollary 1. *Let \mathcal{A} be a set of sensitive attributes. Then \mathcal{A} is also 0-critical.*

Another consequence we can draw is that privacy requirements will always be satisfied if no critical attributes are disseminated.

Corollary 2. *Let \mathcal{O} be a public observations that does not include any critical attributes for a user u and an adversary Adv . Then u 's privacy policy \mathcal{R}_u is σ -satisfied against Adv .*

The corollary above implies that, while we cannot provide general non-disclosure guarantees in open settings, we can provide privacy assessments for specific privacy requirements, given an accurate estimate of the adversary's prior beliefs.

While privacy assessments alone are not satisfactory from a computer security perspective, where we usually require hard security guarantees quantified over all possible adversaries, the fact remains that we are faced with privacy issues in open settings that are to this day unanswered for due to the impossibility of hard guarantees in such settings. Pragmatically thinking, we are convinced that we should move from impossible hard guarantees to more practical privacy assessments instead. This makes particularly sense in settings where users are not victims of targeted attacks, but instead fear attribute disclosure to data-collecting third parties.

5 Linkability in Open Settings

In the following we instantiate the general privacy model introduced in the last section to reason about the likelihood that two profiles of the same user are linked by the adversary in open settings. We introduce the novel notion of (k, d) -anonymity with which we assess anonymity and linkability based on the similarity of profiles within an online community.

To simplify the notation we introduce in this section, we will, in the following, talk about matching *entities* ϵ and ϵ' the adversary wants to link, instead of profiles P_1 and P_2 that belong to the same user u . All definitions introduced in the general framework above naturally carry over to entities as well.

5.1 Model Instantiation for Linkability

In the linkability problem, we are interested in assessing the likelihood that two matching entities ϵ and ϵ' can be linked, potentially across different online platforms. The corresponding privacy requirements, as introduced in Section 4.4, are $r_1 = (\epsilon, \alpha_L)$ and $r_2 = (\epsilon', \alpha_L)$, where α_L is the attribute that ϵ and ϵ' belong to the same user. Consequently, we say that these entities are unlinkable if they satisfy the aforementioned privacy requirements.

Definition 16 (Unlinkability). *Two entities ϵ and ϵ' are σ -unlinkable if $\{\theta_\epsilon, \theta_{\epsilon'}\} \models_\sigma \{r_1, r_2\}$.*

5.2 Anonymity

To assess the identity disclosure risk of an entity ϵ within a collection of entities \mathcal{E} , we use the following intuition: ϵ is anonymous in \mathcal{E} if there is a subset $\mathcal{E}' \subseteq \mathcal{E}$ to which ϵ is very similar. The collection \mathcal{E}' then is an anonymous subset of \mathcal{E} for ϵ .

To assess the similarity of entities within a collection of entities, we will use a distance measure dist on the entity models of these entities. We will require that this measure provides all properties of a metric.

A collection of entities in which the distance of all entities to ϵ is small (i.e., \leq a constant d) is called d -convergent for ϵ .

Definition 17. A collection of entities \mathcal{E} is d -convergent for ϵ if $\text{dist}(\theta_\epsilon, \theta_{\epsilon'}) \leq d$ for all $\epsilon' \in \mathcal{E}$.

Convergence measures the similarity of a collection of individuals. Anonymity is achieved if an entity can find a collection of entities that are all similar to this entity. This leads us to the definition of (k, d) -anonymity, which requires a subset of similar entities of size k .

Definition 18. An entity ϵ is (k, d) -anonymous in a collection of entities \mathcal{E} if there exists a subset of entities $\mathcal{E}' \subseteq \mathcal{E}$ with the properties that $\epsilon \in \mathcal{E}'$, that $|\mathcal{E}'| \geq k$ and that \mathcal{E}' is d -convergent.

An important feature of this anonymity definition is that it provides anonymity guarantees that can be derived from a subset of all available data, but continue to hold once we consider a larger part of the dataset.

Corollary 3. If an entity is (k, d) -anonymous in a collection of entities \mathcal{E} , then it is also (k, d) -anonymous in the collection of entities $\mathcal{E}' \supset \mathcal{E}$.

Intuitively, (k, d) -anonymity is a generalization of the classical notions of k -anonymity to open settings without pre-defined quasi-identifiers. We schematically illustrate such anonymous subsets in Figure 1.

5.3 Entity Matching

We define the notion of *matching* identities. As before, we use the distance measure dist to assess the similarity of two entities.

Definition 19. An entity ϵ c -matches an entity ϵ' if $\text{dist}(\theta_\epsilon, \theta_{\epsilon'}) \leq c$.

Similarly, we can also define the notion of one entity matching a collection of entities.

Definition 20. A collection of entities \mathcal{E} c -matches an entity ϵ' if all entities $\epsilon \in \mathcal{E}$ c -match ϵ' .

Assuming the adversary only has access to the similarity of entities, the best he can do is comparing the distance of all entities $\epsilon \in \mathcal{E}$ to ϵ' and make a probabilistic choice proportional to their relative distance values.




Figure 1: Anonymity in crowdsourcing systems.

Now, if the matching identity ϵ^* is d -convergent in \mathcal{E} the, all entities in \mathcal{E} will have a comparatively similar distance to ϵ' .

Lemma 1. *Let \mathcal{E} be d -convergent for ϵ^* . If ϵ^* c -matches ϵ' , then \mathcal{E} $(c + d)$ -matches ϵ' .*

Proof. Since \mathcal{E} is d -convergent for ϵ^* , $\forall \epsilon' \in \mathcal{E} : \text{dist}(\epsilon^*, \epsilon') \leq d$. Using the triangle inequality, and the fact that ϵ^* c -matches the entity ϵ' , we can bound the distance of all entities $\epsilon \in \mathcal{E}$ to ϵ' by $\forall \epsilon'' \in \mathcal{E} : \text{dist}(\epsilon, \epsilon') \leq c + d$. Hence \mathcal{E} $(c + d)$ -matches the entity ϵ' . \square \square

Hence, the matching entity ϵ^* does not c -match ϵ' for a small value of c , the adversary **Adv** he will have a number of possibly matching entities that are similarly likely to match ϵ' .

We get the same result if not the whole collection \mathcal{E} is convergent, but if there exists a subset of convergent entities that allows the target to remain anonymous.

Corollary 4. *Let ϵ' be (k, d) -anonymous in \mathcal{E} . If ϵ' c -matches an entity ϵ then there is a subset $\mathcal{E}' \subseteq \mathcal{E}$ of size at least k which $(c + d)$ -matches ϵ .*

5.4 Identity Disclosure

We assume that the adversary uses the similarity of the candidate entities to his target entity ϵ' to make his decision. The likelihood that the adversary chooses a specific entity ϵ^* then is the relative magnitude of $\text{dist}(\epsilon^*, \epsilon)$, i.e.

$$Pr[\text{Adv chooses } \epsilon^*] = 1 - \frac{\text{dist}(\epsilon^*, \epsilon')}{\sum_{\epsilon \in \mathcal{E}} \text{dist}(\epsilon, \epsilon')}.$$

We can now bound the likelihood with which a specific entity ϵ^* would be chosen by the adversary if ϵ^* is (k, d) -anonymous.

Theorem 2. *Let the matching entity ϵ^* of the entity ϵ' in the collection $\mathcal{E} = \{\epsilon_1, \dots, \epsilon_n\}$ be (k, d) -anonymous in \mathcal{E} . Furthermore let ϵ^* c -match ϵ' . Then an adversary $\text{Adv} = (b, \emptyset)$ with uniform prior belief b and with empty world knowledge that only observes the similarity of entities links the entity ϵ^* to ϵ' with a likelihood of at most $t \leq 1 - \frac{c}{c+(k-1)(c+d)}$.*

Proof. Let \mathcal{E}^* be the (k, d) anonymous subset of ϵ^* in \mathcal{E} . Let t^* be the likelihood of identifying ϵ^* from \mathcal{E}^* . Then clearly $t < t^*$ since we remove all possible, but wrong candidates in $\mathcal{E} \setminus \mathcal{E}^*$.

Since ϵ^* c -matches ϵ' , by Lemma 1, we can upper bound the distance of each entity in \mathcal{E}^* to ϵ' , i.e.,

$$\forall \epsilon \in \mathcal{E}^* : \text{dist}(\epsilon, \epsilon') \leq c + d$$

We can now bound t^* as follows:

$$\begin{aligned} t^* &= \Pr[\text{Adv chooses } \epsilon] \\ &= 1 - \frac{c}{c + (k-1) \left(\sum_{\epsilon \in \mathcal{E}^* \setminus \{\epsilon^*\}} \text{dist}(\epsilon, \epsilon') \right)} \leq 1 - \frac{c}{c + (k-1)(c+d)} \end{aligned}$$

□

□

Theorem 2 shows that, as long as entities remain anonymous in a suitably large anonymous subset of a collection of entities, an adversary will have difficulty identifying them with high likelihood. Recalling our unlinkability definition from the beginning of the section, this result also implies that ϵ^* is σ -unlinkable for $\sigma = t$.

Corollary 5. *Let the matching entity ϵ^* of the entity ϵ' in the collection $\mathcal{E} = \{\epsilon_1, \dots, \epsilon_n\}$ be (k, d) -anonymous in \mathcal{E} . Then ϵ^* and ϵ' are σ -unlinkable for $\sigma = 1 - \frac{c}{c+(k-1)(c+d)}$ against an adversary $\text{Adv} = (b, \emptyset)$ with uniform prior belief and empty world knowledge that only observes entity similarity.*

In Section 6.5 we present experiments that evaluate the anonymity and linkability of individuals in the Online Social Network Reddit, and measure how well they can be identified from among their peers.

5.5 Limitations

The quality of the assessment provided by the d -convergence model largely depends on adversarial prior belief we assume: in our results above, we assume an adversary without any prior knowledge. In practice, however, the adversary might have access to prior beliefs that can help him in his decision making. Therefore, turning such assessments into meaningful estimates in practice requires a careful estimation of prior knowledge by, e.g., producing a more accurate profile model: the problem of comprehensive profile building for entities in an open setting is an open question that has been examined somewhat in the literature [5, 9, 8, 28, 4], but on the whole still leaves a lot of space for future work.

This concludes the formal definitions of our d -convergence model. In the next sections, we instantiate it for identity disclosure risk analyses based on user-generated text-content and apply this instantiation to the OSN Reddit.

6 Linkability Evaluation on Reddit

While the main focus of this paper is to present the actual privacy model as such, the following experiments are meant to provide first insights into the application of our framework, without taking overly complex adversarial capabilities into account. The evaluation can easily be extended to a more refined model of an adversary without conceptual difficulties.

We first articulate the goals of this evaluation, and then, secondly, describe the data collection process, followed by defining the instantiation of the general framework we use for our evaluation in the third step. Fourth, we introduce the necessary processing steps on our dataset, before we finally discuss the results of our evaluation.

6.1 Goals

In our evaluation, we aim at validating our model by conducting two basic experiments. First, we want to empirically show that, our model instantiation yields a suitable abstraction of real users for reasoning about their privacy. To this end, profiles of the same user should be more similar to each other (less distant) than profiles from different users.

Second, we want to empirically show that a larger anonymous subset makes it more difficult for an adversary to correctly link the profile. Thereby, we inspect whether anonymous subsets provide a practical estimate of a profile’s anonymity.

Given profiles with anonymous subsets of similar size, we determine the percentage of profiles which the adversary can match within the top k results, i.e., given a source profile, the adversary computes the top k most similar (less distant) profiles in the other subreddit. We denote this percentage by *precision@k* and correlate it to the size of the anonymous subsets.

We fix the convergence of the anonymous subsets to be equal to the matching distance between two corresponding profiles. Our intuition is that, this way, the anonymous subset captures most of the profiles an adversary could potentially consider matching.

6.2 Data-Collection

For the empirical evaluation of our privacy model, we use the online social network Reddit [1] that was founded in 2005 and constitutes one of the largest discussion and information sharing platforms in use today. On Reddit, users share and discuss topics in a vast array of topical subreddits that collect all

topics belonging to one general area; e.g. there are subreddits for world news, tv series, sports, food, gaming and many others. Each subreddit contains so-called submissions, i.e., user-generated content that can be commented on by other users.

To have a ground truth for our evaluation, we require profiles of the same user same user across different OSNs to be linked. Fortunately, Reddit’s structure provides an inherent mechanism to deal with this requirement. Instead of considering Reddit as a single OSN, we treat each subreddit as its own OSN. Since users are identified through the same pseudonym in all of those subreddits, they remain linkable across the subreddits’ boundaries. Hence our analysis has the required ground truth. The adversary we simulate, however, is only provided with the information available in the context of each subreddit and thus can only try to match profiles across subreddits. Ground truth in the end allows us to verify the correctness of his match.

To build up our dataset, we built a crawler using Reddit’s API to collect comments. Recall that subreddits contain submissions that, in turn, are commented by the users. For our crawler, we focused on the large amount of comments because they contain a lot of text and thus are best suitable for computing the unigram models.

Our crawler operates in two steps that are repeatedly executed over time. During the whole crawling process, it maintains a list of already processed users. In the first step, our crawler collects a list of the overall newest comments on Reddit from Reddit’s API and inserts these comments into our dataset. In the second step, for each author of these comments who has not been processed yet, the crawler also collects and inserts her latest 1,000 comments into our dataset. Then, it updates the list of processed users. The number of 1,000 comments per user, is a restriction of Reddit’s API.

In total, during the whole September 2014, we collected more than 40 million comments from over 44,000 subreddits. The comments were written by about 81,000 different users which results in more than 2.75 million different profiles.

The whole dataset is stored in an anonymized form in a MySQL database and is available upon request.

6.3 Model Instantiation

On Reddit, users only interact with each other by by posting comments to text of link submissions. Reddit therefore does not allow us to exploit features found in other social networks, such as friend links or other static data about each user. On the other hand, this provides us with the opportunity to evaluate the linkability model introduced in Section 5 based dynamic, user-generated content, in this case user-generated text content.

Since we only consider text content, we instantiate the general model from the previous sections with an unigram model, where each attribute is a word unigram, an its value is the frequency with which the unigram appears in the profiles comments. Such unigram models have succesfully been used in the past

to characterize the information within text content and to correlate users across different online platforms [15, 23].

Definition 21 (Unigram Model). *Let \mathcal{V} be a finite vocabulary. The unigram model $\theta_P = p_i$ of a profile is a set of frequencies $p_i \in [0, \dots, 1]$ with which each unigram $w_i \in \mathcal{V}$ appears in the profile P . Each frequency p_i is determined by*

$$p_i = \frac{\text{count}(w_i, P)}{\sum_{w \in \mathcal{V}} \text{count}(w, P)}$$

Since the unigram model essentially constitutes a probability distribution, we instantiate our distance metric `dist` with the Jensen-Shannon divergence [14]. The Jensen-Shannon divergence is a symmetric extension of the Kullback-Leiber divergence has been shown to be successful in many related information retrieval scenarios.

Definition 22. *Let P and Q be two statistical models over a discrete space Ω . The Jensen-Shannon divergence is defined by*

$$D_{\text{JS}} = \frac{1}{2} D_{\text{KL}}(P||M) + \frac{1}{2} D_{\text{KL}}(Q||M)$$

where D_{KL} is the Kullback-Leibler divergence

$$D_{\text{KL}}(P||Q) = \sum_{\omega \in \Omega} \log \left(\frac{P(\omega)}{Q(\omega)} \right) P(\omega)$$

and M is the averaged distribution $M = \frac{1}{2}(P + Q)$.

In the following, we will use the square-root of the Jensen-Shannon divergence, constituting a metric, as our distance measure, i.e., $\text{dist} = \sqrt{D_{\text{JS}}}$.

6.4 Data-Processing

The evaluation on our dataset is divided into sequentially performed computation steps, which include the normalization of all comments, the computation of unigram models for each profile, a filtering of our dataset to keep the evaluation tractable, the computation of profile distances and the computation of (k, d) -anonymous subsets.

Normalizing Comments. Unstructured, heterogeneous data, as in our case, may contain a variety of valuable information about a user’s behavior, e.g., including formatting and punctuation. Although we could transform these into attributes, we do not consider them here for the sake of simplicity.

In order to get a clean representation to apply the unigram model on, we apply various normalization steps, including transformation to lower case, the removal of Reddit formatting and punctuation except for smilies. Moreover, we apply a encoding specific normalization, replace URLs by their hostnames and shorten repeated characters in words like `coooo1` to a maximum of three.

Finally, we also filter out a list of 597 stopwords from the comments. Therefore, we perform six different preprocessing steps on the data, which we describe in more detail in the following.

1. **Convert to lower case letters:** In our statistical language models, we do not want to differentiate between capitalized and lowercased occurrences of words. Therefore, we convert the whole comment into lower case.
2. **Remove Reddit formatting:** Reddit allows users to use a wide range of formatting modifiers that we divide into two basic categories: formatting modifiers that influence the typography and the layout of the comment, and formatting modifiers that include external resources into a comment. The first kind of modifier, named layout modifiers, is stripped off the comment, while leaving the plain text. The second kind of modifier, called embedding modifiers, is removed from the comment completely.

One example for a layout modifier is the asterisk: When placing an asterisk both in front and behind some text, e.g., `*text*`, this text will be displayed in italics, e.g., *text*. Our implementation removes these enclosing asterisks, because they are not valuable for computing statistical language models for n -grams and only affect the layout. Similarly, we also remove other layout modifiers such as table layouts, list layouts and URL formatting in a way that only the important information remains.

A simple example for embedding modifiers are inline code blocks: Users can embed arbitrary code snippets into their comments using the ‘`code`’ modifier. Since these code blocks do not belong to the natural language part of the comment and only embed a kind of external resource, we remove them completely. In addition to code blocks, the category of embedding modifiers also includes quotes of other comments.

3. **Remove stacked diacritics:** In our dataset, we have seen that diacritics are often misused. Since Reddit uses Unicode as its character encoding, users can create their own characters by arbitrarily stacking diacritics on top of them. To avoid this kind of unwanted characters, we first normalize the comment by utilizing the unicode character composition, which tries to combine each letter and its diacritics into a single precombined character. Secondly, we remove all remaining diacritic symbols from the comment. While this process preserves most of the normal use of diacritics, it is able to remove all unwanted diacritics.
4. **Replace URLs by their hostname:** Generally, a URL is very specific and a user often does not include the exact same URL in different comments. However, it is much more common that a user includes different URLs that all belong to the same hostname, e.g., `www.mypage.com`. Since our statistical language models should represent the expected behavior of a user in terms of used words (including URLs), we restrict all URLs to their hostnames.

Top	Reddit		subreddit: Lost		subreddit: TipOfMyTongue	
	Unigram	Frequency	Unigram	Frequency	Unigram	Frequency
1.	people	4,127,820	island	832	www.youtube.com	3663
2.	time	2,814,841	show	750	song	1,542
3.	good	2,710,665	lost	653	remember	1,261
4.	gt	2,444,240	time	580	en.wikipedia.org	1,100
5.	game	1,958,850	people	527	sounds	1,007
6.	pretty	1,422,640	locke	494	solved	924
7.	2	1,413,118	season	431	movie	918
8.	lot	1,385,167	jacob	429	find	829
9.	work	1,352,292	mib	372	:)	786
10.	1	1,184,029	jack	310	game	725
11.	3	1,124,503	episode	280	time	678
12.	great	1,070,299	ben	255	thinking	633
13.	point	1,063,239	good	250	good	633
14.	play	1,060,985	monster	237	www.imdb.com	584
15.	years	1,032,270	lot	220	video	583
16.	bad	1,008,607	gt	182	pretty	570
17.	day	989,180	character	165	youtu.be	569
18.	love	988,567	walt	163	mark	548
19.	find	987,171	man	162	edit	540
20.	shit	976,928	dharma	162	post	519

Table 1: Top 20 unigrams of Reddit and two sample subreddits Lost and TipOfMyTongue.

5. **Remove punctuation:** Most of the punctuation belongs to the sentence structure and, thus, should not a part of our statistical language models. Therefore, we remove all punctuation except for the punctuation inside URLs and smilies. We do not remove the smilies, because people are using them in a similar role as words to enrich their sentences: Every person has her own subset of smilies that she typically uses. To keep the smilies in the comment, we maintain a list of 153 different smilies that will not be removed from the comment.
6. **Remove duplicated characters:** In the internet, people often duplicate characters in a word to add emotional nuances to their writing, e.g., coooooooooo1. But sometimes the number of reduplicated characters varies, even if the same emotion should be expressed. Thus, we reduce the number of duplicated characters to a maximum of 3, e.g., cooo1. In practice, this truncation allows us to differentiate between the standard use of a word and the emotional variation of it, while it does not depend on the actual number of duplicated characters.

Computing Unigram Models. From the normalized data, we compute

the unigram frequencies for each comment. Recall that our dataset consists of many subreddits that each form their own OSN. Thus, we aggregate the corresponding unigram frequencies per profile, per subreddit, and for Reddit as a whole. Using this data, we compute the word unigram frequencies for each comment as described in Section 6.3.

Since a subreddit collects submissions and comments to a single topic, we expect the unigrams to reflect its topic specific language. Indeed, the 20 most frequently used unigrams of a subreddit demonstrate that the language adapts to the topic. As an example, we show the top 20 unigrams (excluding stopwords) of Reddit and two sample subreddits *Lost* and *TipOfMyTongue* in Table 1. As expected, there are subreddit specific unigrams that occur more often in the context of one subreddit than in the context of any other subreddit. For example, the subreddit *Lost* deals with a TV series that is about the survivors of a plane crash and its aftermath on an island. Unsurprisingly, the word *island* is the top unigram in this subreddit. In contrast, the subreddit *TipOfMyTongue* deals with the failure to remember a word from memory and, thus, has the word *remember* in the list of its top three unigrams.

Filtering the Dataset. To reduce the required amount of computations we restrict ourselves to *interesting profiles*. We define an interesting profile as one that contains at least 100 comments and that belongs to a subreddit with at least 100 profiles. Additionally, we dropped the three largest subreddits from our dataset to speed up the computation.

In conclusion, this filtering results in 58,091 different profiles that belong to 37,935 different users in 1,930 different subreddits.

Distances Within and Across Subreddits. Next, we compute the pairwise distance within and across subreddits using our model instantiation. Excluding the distance of profiles to themselves, the minimal, maximal and average distance of two profiles within subreddits in our dataset are approximately 0.12, 1 and 0.79 respectively. Across subreddits, the minimal, maximal and average distance of two profiles are approximately 0.1, 1 and 0.85 respectively.

Anonymous Subsets. Utilizing the distances within subreddits, we can determine the anonymous subsets for each profile in a subreddit. More precisely, we compute the anonymous subset for each pair of profiles from the same user. We set the convergence d to the matching distance between both profiles and determine the size of the resulting anonymous subset.

6.5 Evaluation and Discussion

In this subsection, we inspect and interpret the results of our experiments with regard to our aforementioned goals. Therefore, we first start by giving evidence that our approach indeed provides a suitable abstraction of real users for reasoning about their privacy.

To this end, we compare the distance of matching profiles to the average distance of non-matching profiles. In particular, for each pair of profiles from the same user in subreddits s and s' , we plot the average distance from the




Figure 2: The average distance between a profile in subreddit s and all profiles in s' versus the matching distance between the profile and its correspondence in s' .




Figure 3: The anonymous subset size correlated to the precision an adversary has if considering the top 5 profiles as matching.

profile in s to the non-matching profiles in s' in relation to the distance to the matching profile in s' in Figure 2. The red line denotes the function $y = x$ and divides the figure into two parts: if a point lies below the line through the origin, the corresponding profiles match better than the average of the remaining profiles. Since the vast majority of datapoints is located below the line, we can conclude that profiles of the same user match better than profiles of different users.

Our second goal aimed at showing that anonymous subsets indeed can be used to reason about the users' privacy. Therefore, we investigate the chances of an adversary to find a profile of the same user within the top k matches and

relate its chance to the size of the profile’s anonymous subset. More precisely, given multiple target profiles with similar anonymous subset sizes, we determine the, so called, $\text{precision@}k$, i.e., the ratio of target profiles that occur in the top k ranked matches (by ascending distance from the source profiles). We relate this $\text{precision@}k$ to the anonymous subset sizes with a convergence d set to the distance between the source and target profiles, and we group the anonymous subset sizes in bins of size 10.

In our evaluation, we considered $k \in \{1, 5, 10, 20\}$, which all yield very similar results. Exemplarily, we correlate the aforementioned measures for $k = 5$ in Figure 3, clearly showing that an increasing anonymous subset size correlates with an increasing uncertainty – i.e., decreasing precision – for the adversary.

7 Conclusion and Future Work

We presented a user-centric privacy framework for reasoning about privacy in open web settings. In our formalization, we address the essential challenges of privacy in open settings: we defined a comprehensive data model that can deal with the unstructured dissemination of heterogeneous information, and we derived the sensitivity of information from user-specified and context-sensitive privacy requirements. We showed that, in this formalization of privacy in open settings, hard security guarantees in the sense of Differential Privacy are impossible to achieve. We then instantiated the general framework to reason about the identity disclosure problem. The technical core of our identity disclosure model is the new notion of (k, d) -anonymity that assesses the anonymity of entities based on their similarity to other entities within the same community. We applied this instantiation to a dataset of 15 million user-generated text entries collected from the Online Social Network Reddit and showed that our framework is suited for the assessment of linkability threats in Online Social Networks.

As far as future work is concerned, many directions are highly promising. First, our general framework only provides a static view on privacy in open settings. Information dissemination on the Internet, however, is, in particular, characterized by its highly dynamic nature. Extending the model presented in this paper with a suitable transition system to capture user actions might lead to powerful system for monitoring privacy risks in dynamically changing, open settings. Second, information presented in Online Social Networks is often highly time-sensitive, e.g., shared information is often only valid for a certain period of time, and personal facts can change over time. Explicitly including timing information in our entity model will hence further increase the accuracy of the entity models derived from empirical evidence. Finally, our privacy model is well-suited for the evaluation of protection mechanisms for very specific privacy requirements, and new such mechanisms with provable guarantees against restricted adversaries can be developed. On the long run, we pursue the vision of providing the formal foundations for comprehensive, trustworthy privacy assessments and, ultimately, for developing user-friendly privacy assessment tools.

References

- [1] The online social network reddit. <http://www.reddit.com>. Accessed Sep 2015.
- [2] Directive 95/46/EC of the European Parliament and of the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data, 1996.
- [3] Michael Backes, Aniket Kate, Praveen Manoharan, Sebastian Meiser, and Esfandiar Mohammadi. AnoA: A Framework for Analyzing Anonymous Communication Protocols. In *Proceedings of the 26th IEEE Computer Security Foundations Symposium (CSF)*, pages 163–178, 2013.
- [4] Marco Balduzzi, Christian Platzer, Thorsten Holz, Engin Kirda, Davide Balzarotti, and Christopher Kruegel. Abusing Social Networks for Automated User Profiling. In *Proceedings of the 13th international Conference on Recent Advances in Intrusion Detection (RAID)*, pages 422–441, 2010.
- [5] Andrea Calì, Diego Calvanese, Simona Colucci, Tommaso Di Noia, and FrancescoM. Donini. A Logic-Based Approach for Matching User Profiles. In *Knowledge-Based Intelligent Information and Engineering Systems*, pages 187–195. 2004.
- [6] Konstantinos Chatzikokolakis, MiguelE. Andrés, NicolásEmilio Bordenabe, and Catuscia Palamidessi. Broadening the Scope of Differential Privacy Using Metrics. In *Proceedings of the 13th Privacy Enhancing Technologies Symposium (PETS)*, pages 82–102. 2013.
- [7] Rui Chen, Benjamin C.M. Fung, Philip S. Yu, and Bipin C. Desai. Correlated network data publication via differential privacy. *The VLDB Journal*, 23(4):653–676, 2014.
- [8] Terence Chen, Mohamed Ali Kaafar, Arik Friedman, and Roksana Boreli. Is More Always Merrier?: A Deep Dive into Online Social Footprints. In *Proceedings of the 2012 ACM Workshop on Online Social Networks (WOSN)*, pages 67–72, 2012.
- [9] Keith Cortis, Simon Scerri, Ismael Rivera, and Siegfried Handschuh. Discovering Semantic Equivalence of People Behind Online Profiles. In *Proceedings of the 5th International Workshop on Resource Discovery (RED)*, pages 104–118, 2012.
- [10] Irit Dinur and Kobbi Nissim. Revealing Information While Preserving Privacy. In *Proceedings of the 22nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 202–210, 2003.

- [11] Cynthia Dwork. Differential Privacy: A Survey of Results. In *Proceedings of the 5th International Conference on Theory and Applications of Models of Computation*, pages 1–19, 2008.
- [12] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *Advances in Cryptology (EUROCRYPT)*, page 486–503, 2006.
- [13] Cynthia Dwork and Moni Naor. On the difficulties of disclosure prevention in statistical databases or the case for differential privacy. *Journal of Privacy and Confidentiality*, 2(1):8, 2008.
- [14] Dominik M. Endres and Johannes E. Schindelin. A new metric for probability distributions. *IEEE Transactions on Information Theory*, 49(7):1858–1860, 2003.
- [15] Oana Goga, Howard Lei, SHK. Parthasarathi, Gerald Friedland, Robin Sommer, and Renata Teixeira. Exploiting innocuous activity for correlating users across sites. In *WWW*, 2013.
- [16] Raymond Heatherly, Murat Kantarcioglu, and Bhavani Thuraisingham. Preventing Private Information Inference Attacks on Social Networks. *IEEE Transactions on Knowledge and Data Engineering*, 25(8):1849–1862, 2013.
- [17] Shiva Prasad Kasiviswanathan and Adam Smith. On the ‘Semantics’ of Differential Privacy: A Bayesian Formulation. *Journal of Privacy and Confidentiality*, 6(1):1–16, 2014.
- [18] Daniel Kifer and Ashwin Machanavajjhala. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data*, pages 193–204, 2011.
- [19] Michal Kosinski, Yoram Bachrach, Pushmeet Kohli, David Stillwell, and Thore Graepel. Manifestations of User Personality in Website Choice and Behaviour on Online Social Networks. *Machine Learning*, 95(3):357–380, 2014.
- [20] Victor Lavrenko and W. Bruce Croft. Relevance based language models. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 120–127, 2001.
- [21] Ninghui Li and Tiancheng Li. t-closeness: Privacy beyond k-anonymity and l-diversity. In *Proceedings of the 23rd International Conference on Data Engineering (ICDE)*, 2007.
- [22] Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkitasubramaniam. L-Diversity: Privacy Beyond K-anonymity. *ACM Trans. Knowl. Discov. Data*, 1(1), 2007.

- [23] Mishari Al Mishari and Gene Tsudik. Exploring linkability of user reviews. In *ESORICS*, 2012.
- [24] Arvind Narayanan and Vitaly Shmatikov. Myths and Fallacies of "Personally Identifiable Information". *Communications of the ACM*, 53(6):24–26.
- [25] Arvind Narayanan and Vitaly Shmatikov. De-anonymizing Social Networks. In *Proceedings of the 30th IEEE Symposium on Security and Privacy (S&P)*, pages 173–187, 2009.
- [26] Jay M. Ponte and W. Bruce Croft. A Language Modeling Approach to Information Retrieval. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 275–281, 1998.
- [27] Ronald Rosenfeld. Two decades of statistical language modeling: where do we go from here? *Proceedings of the IEEE*, 88(8):1270–1278, 2000.
- [28] Naveen Kumar Sharma, Saptarshi Ghosh, Fabricio Benevenuto, Niloy Ganguly, and Krishna Gummadi. Inferring Who-is-who in the Twitter Social Network. In *Proceedings of the 2012 ACM Workshop on Workshop on Online Social Networks (WSOON)*, pages 55–60, 2012.
- [29] Fei Song and W. Bruce Croft. A general language model for information retrieval. In *Proceedings of the Eighth International Conference on Information and Knowledge Management (CIKM)*, pages 316–321, 1999.
- [30] Latanya Sweeney. K-Anonymity: A Model for Protecting Privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):557–570, 2002.
- [31] Özlem Uzuner and Boris Katz. A Comparative Study of Language Models for Book and Author Recognition. In *Natural Language Processing – IJCNLP 2005*, pages 969–980. 2005.
- [32] ChengXiang Zhai. Statistical language models for information retrieval a critical review. *Found. Trends Inf. Retr.*, 2(3):137–213, 2008.
- [33] Chengxiang Zhai and John Lafferty. A study of smoothing methods for language models applied to information retrieval. *ACM Transactions on Information Systems (TOIS)*, 22(2):179–214, 2004.
- [34] Elena Zheleva and Lise Getoor. To Join or Not to Join: The Illusion of Privacy in Social Networks with Mixed Public and Private User Profiles. In *Proceedings of the 18th International Conference on World Wide Web (WWW)*, pages 531–540, 2009.
- [35] Elena Zheleva and Lise Getoor. Privacy in social networks: A survey. In *Social Network Data Analytics*, pages 277–306. 2011.

- [36] Bin Zhou and Jian Pei. The k-anonymity and l-diversity approaches for privacy preservation in social networks against neighborhood attacks. *Knowledge and Information Systems*, 28(1):47–77, 2011.